Online Systems Tutorial

11-May-2004

S. Fuess

Contents

- Online functions
- Network view
- Event data flow
- Node, disk assignments
- Accounts
- Access controls
- Monitoring
- Web servers
- Control Room consoles

Online Functions

- Infrastructure
  - Network
  - Storage systems
    - Network file system (NFS)
  - User information
    - NIS / yp
  - Backup system
    - BRU
- Control Room
  - User environment
- Monitoring
  - Examine platforms
- Event Data path
  - From Level 3 to FCC
  - Including event metadata
- Control System
  - Hardware monitoring
    - EPICS
  - “SDAQ”
    - Alternative readout path
- ORACLE Database
  - Primarily “pass through”, with data propagated to Offline database

Online network view

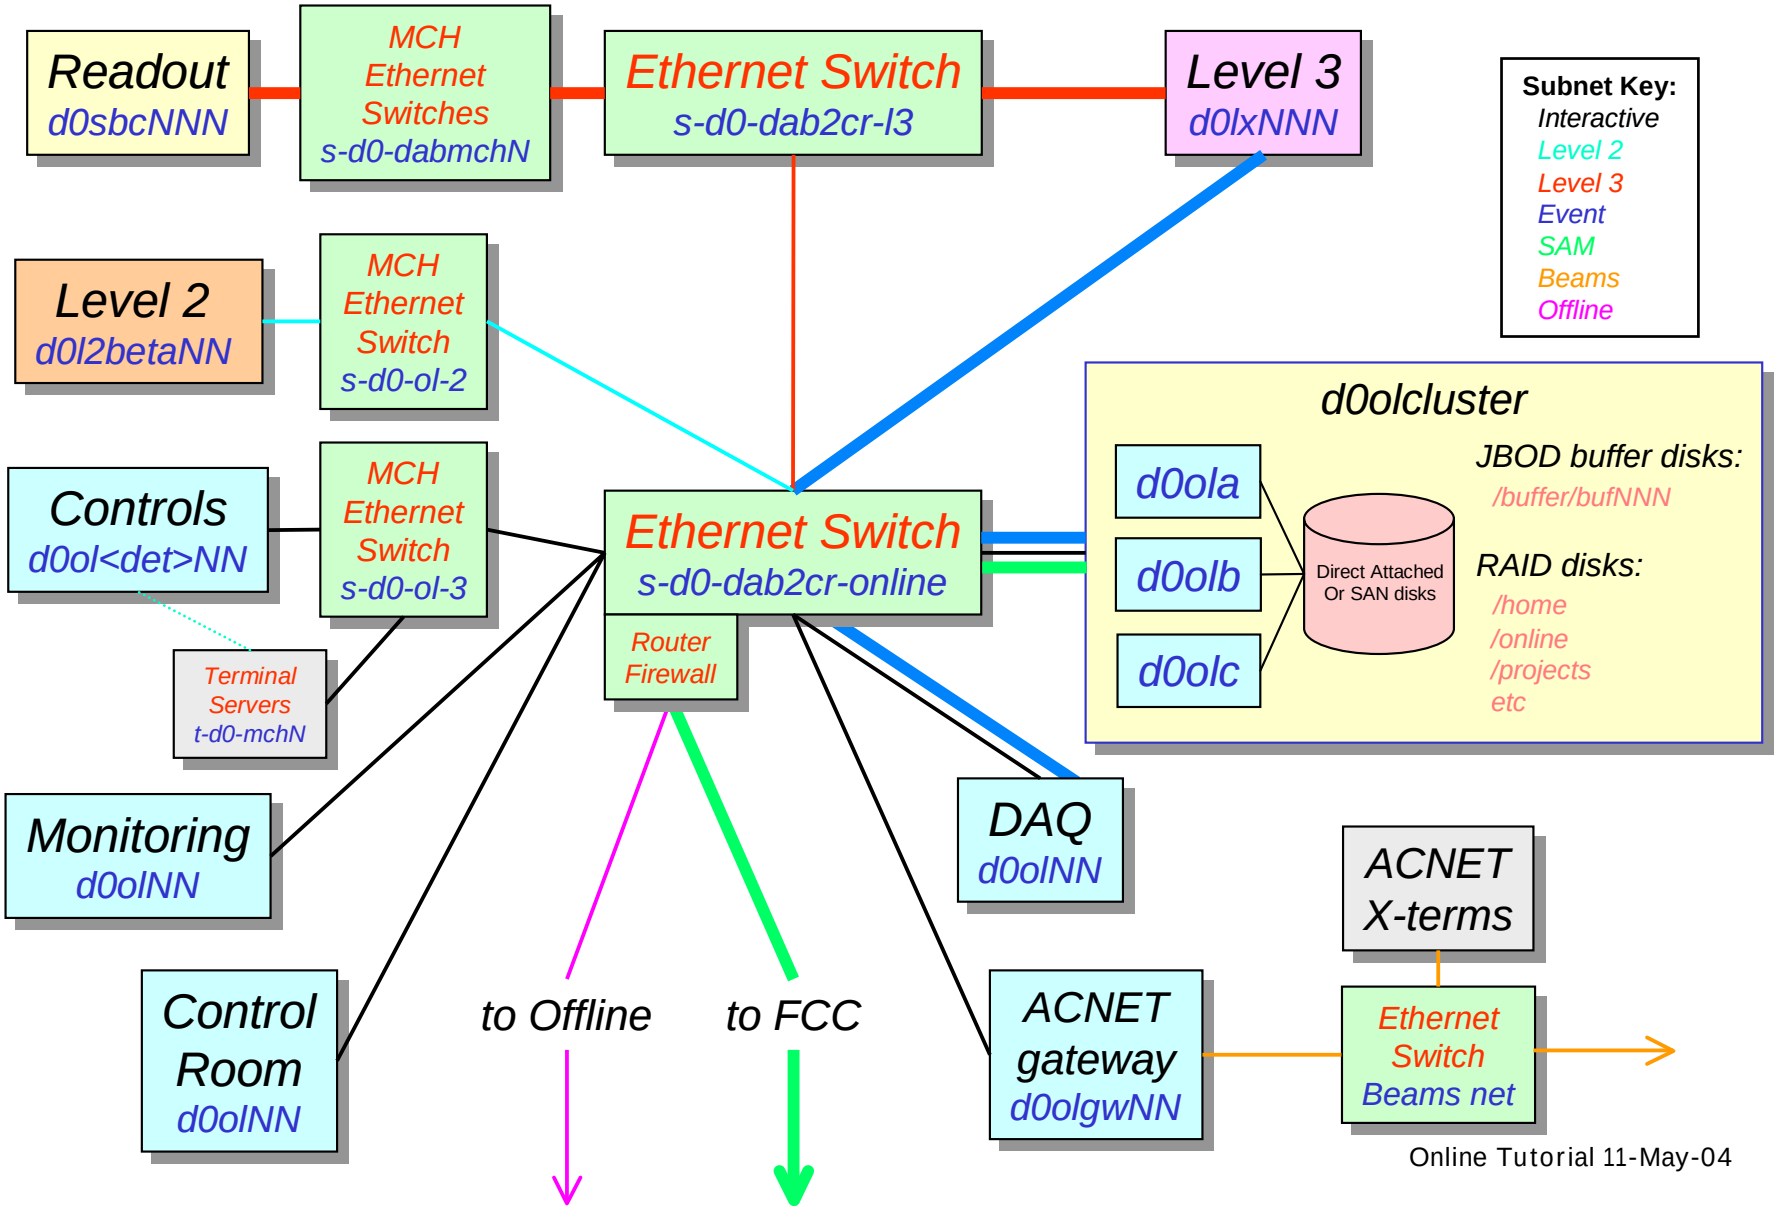

Event Data Flow

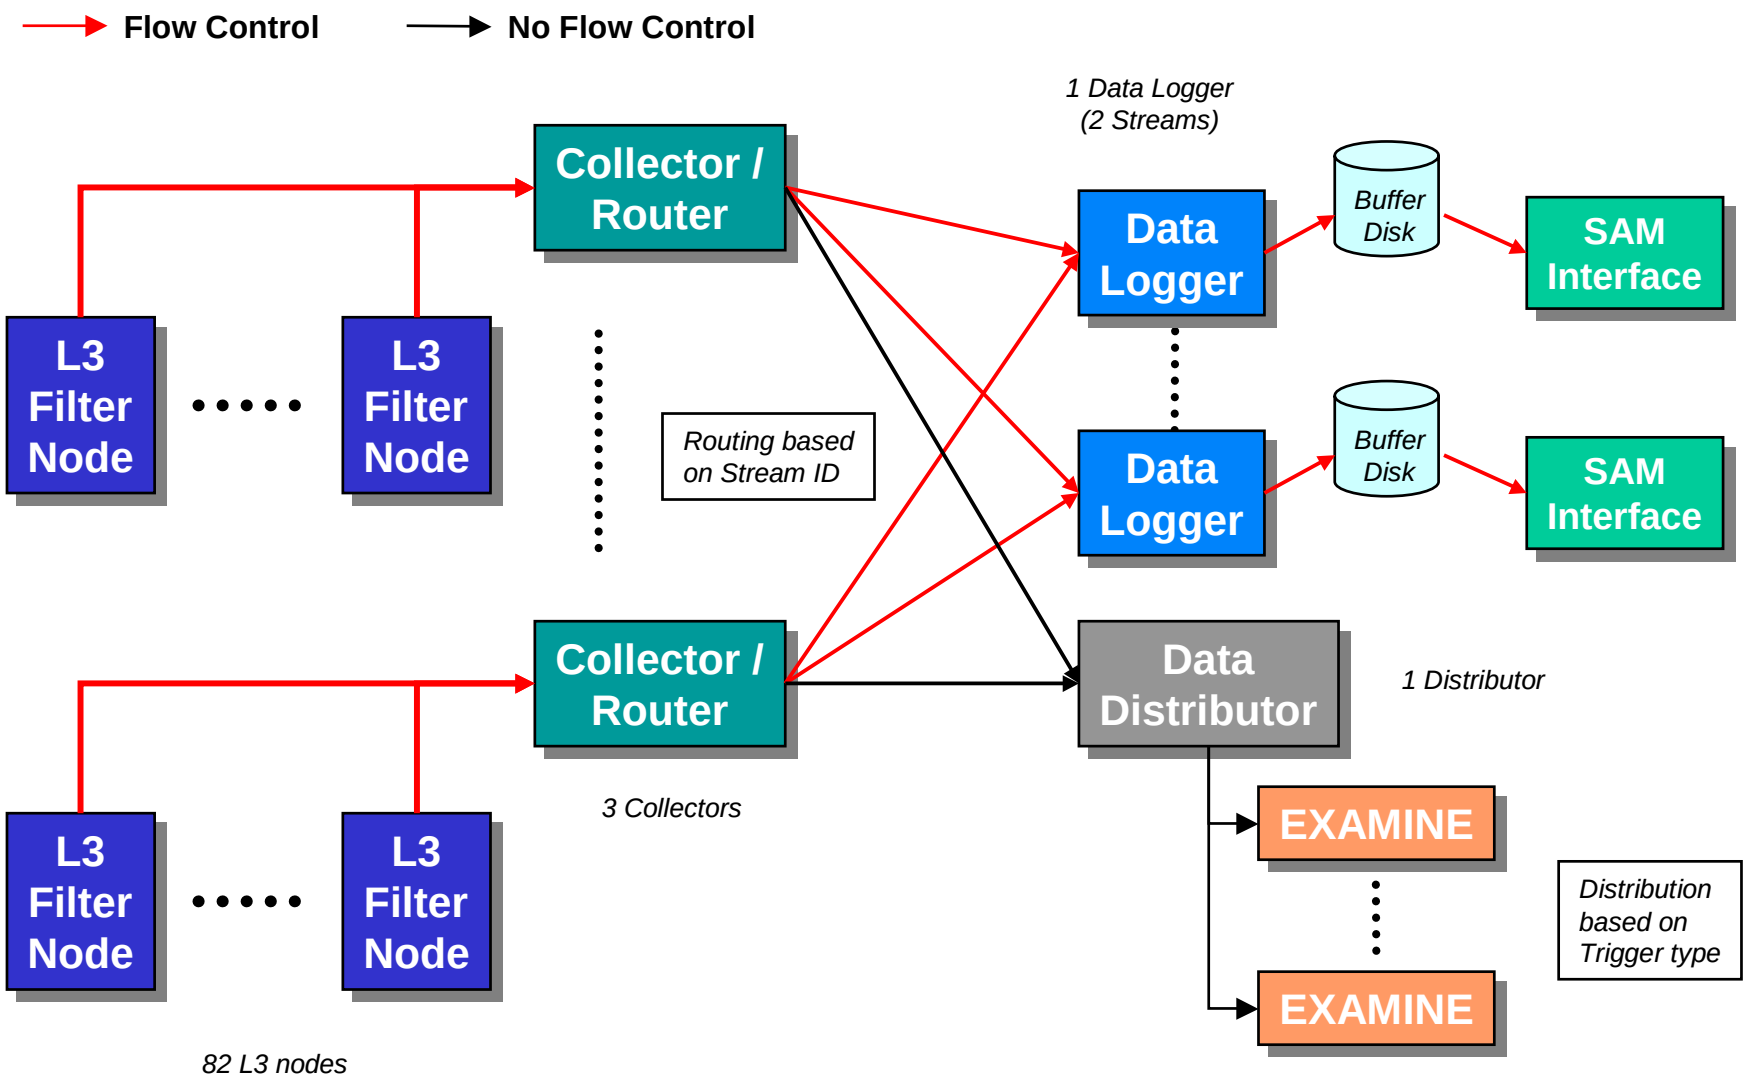

Assignments

- Application / service assignments kept at

  /online/data/d0online_names/d0online_names.py

  If a node dies and the application is relocated, this file must be edited and the instructions within followed

- Node assignments

  http://www-d0online.fnal.gov/www/sys/operations/node_assignments.txt

  http://www-d0online.fnal.gov/www/sys/operations/group_assignments.txt

- Disk assignments

  http://www-d0online.fnal.gov/www/sys/operations/disk_assignments.txt

Accounts

- Two important factors:
  - Authorization – that an account is present for a user on a node
    - Granted to any DO user with need
      - Access to group account may be sufficient
      - Check with “ypcat passwd” or “ls /home”
        - A null /home/<user>/.k5login file indicates the account is locked out!
  - Authentication – that one can demonstrate knowledge of a password
    - The only allowed mode of access originating from outside of the Online system is by Kerberos
      - Almost… there are some Windows nodes with very restricted external access

Accounts

- On the “interactive” (Control Room, Monitoring, Host) systems
  - Authorization
    - Local accounts (e.g. root) for system use only
    - NIS accounts for personal and group users
      - NIS domain server is d0olcluster
      - Personal accounts are “locked out” from non-Kerberos authentication
  - Authentication
    - Only root account has a local password
      - Kerberos .k5login access for remote logins
        - Personal Kerberos credentials (i.e. user@FNAL.GOV)
    - Group NIS accounts
      - NIS password only for local logins
      - Kerberos .k5login access for remote logins
        - Personal Kerberos credentials (i.e. user@FNAL.GOV)
        - Keytab Kerberos credentials (i.e. d0cap/d0/d0ol04.fnal.gov@FNAL.GOV)
    - Personal NIS accounts
      - Kerberos password for local logins (on most nodes)
      - Kerberos or .k5login access for remote logins
        - If a .k5login exists, then must include own credentials

Accounts

- On the “DAQ” (Readout, Level 2, Level 3) systems
  - Authorization
    - Local accounts for system, DAQ, and expert users
  - Authentication
    - Only root account has (should have) a local password
      - Kerberos .k5login access for remote logins
        - Personal Kerberos credentials (i.e. user@FNAL.GOV)
    - DAQ local accounts
      - Kerberos .k5login access for remote logins
        - Personal Kerberos credentials (i.e. user@FNAL.GOV)
        - Keytab Kerberos credentials (i.e. d0run/d0/d0ol07.fnal.gov@FNAL.GOV)
    - Expert user local accounts
      - Kerberos or .k5login access for remote logins
        - Personal Kerberos credentials (i.e. user@FNAL.GOV)

Accounts

- On the Controls systems
  - Authorization
    - Local accounts for expert users
  - Authentication
    - Expert user local accounts
      - Local password for local login
    - No Kerberos! Remote logins are not allowed, and blocked by Online router

Accounts

- Some useful commands
  - To check group account access, e.g.

    cat /home/d0cap/.k5login

  - To see if a user has an NIS account, e.g.

    ypcat passwd | grep fuess

  - To remotely log in to group account on an Online node, e.g.

    kinit fuess

    ssh –l d0cap d0ol04

  - To log in to another node from a group account, e.g. as d0run

    setup d0online

    d0ssh –l d0cap d0ol04

Accounts

- Kerberos keytab files

  How does this work?

  setup d0online

  d0ssh –l d0cap d0ol04

  On each node there is a specific (perhaps empty) set of files of the sort

  /var/adm/krb5/d0smt_keytab

  accessible only by the named user

  These contain the Kerberos key which allows the specific group (e.g. d0smt) account on that node (e.g. d0ol44) to obtain a principal which is of the form

  d0smt/d0/d0ol44.fnal.gov@FNAL.gov

  This principal is then listed within the .k5login of any account for which access is needed

  The d0ssh script is only 2 lines: kinit and ssh

Access controls

- Essential components of the computer security plan for the Online system are that:
  - The detector can operate with the Online system completely isolated from the external world
    - Well-defined isolation points
    - Can isolate from Offline, FCC, or both
    - Local versions of essential services
      - DNS server
      - KDC
    - Sufficient space to buffer event data for > 24 hours
  - Network access to the Online system is tightly controlled
    - Enforced by router module in Online switch acting as a “firewall”
    - Policy is “default deny”

All this leads to functional limitations and operational confusion…

Monitoring – Big Brother

Big Brother main display
click buttons for more info

http://www-d0online/bb

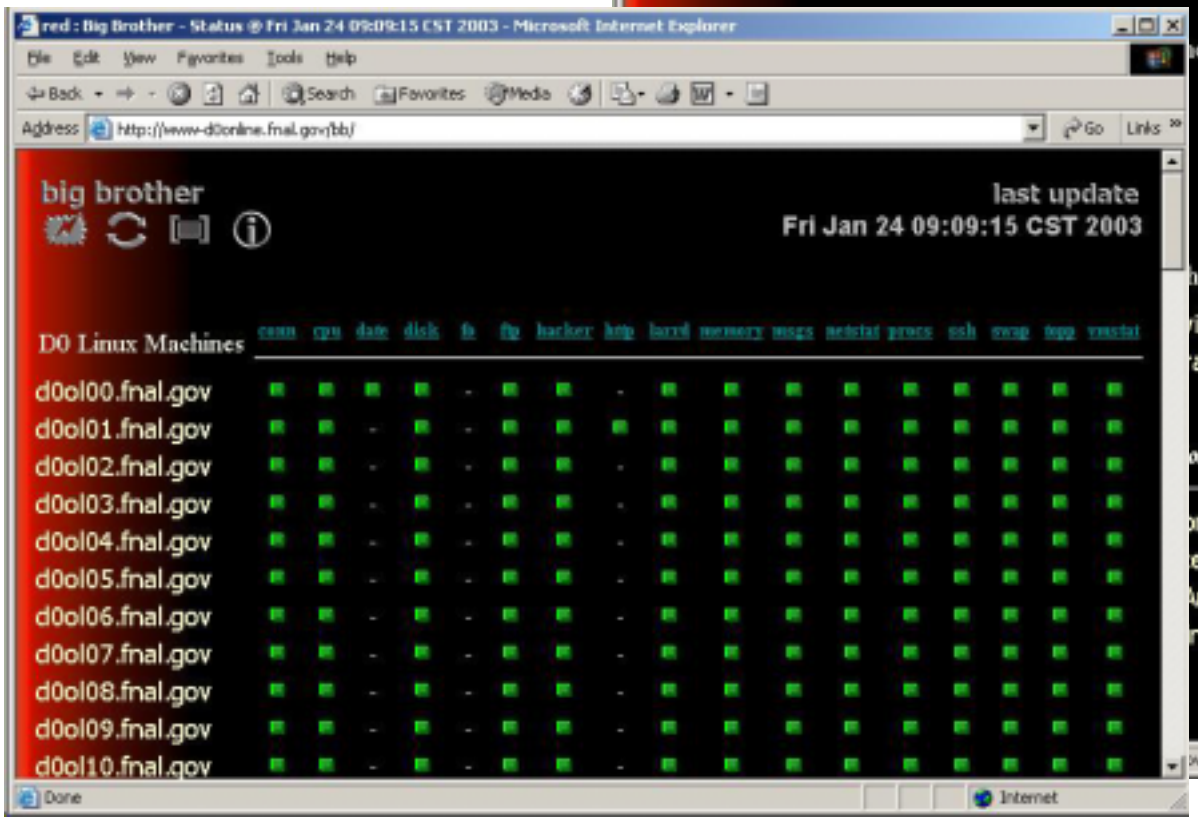

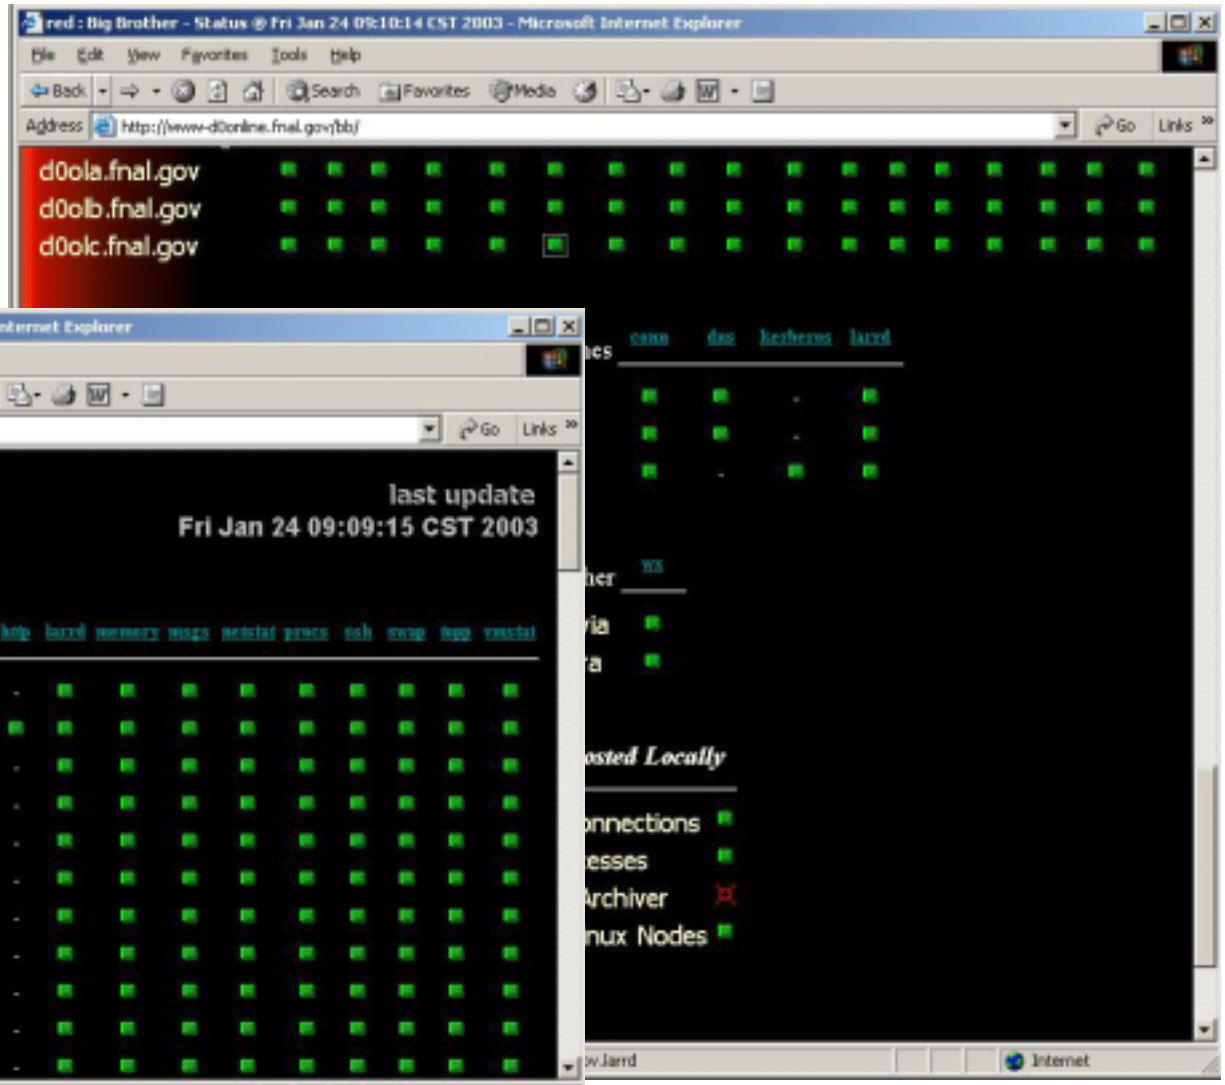

Monitoring – Big Brother

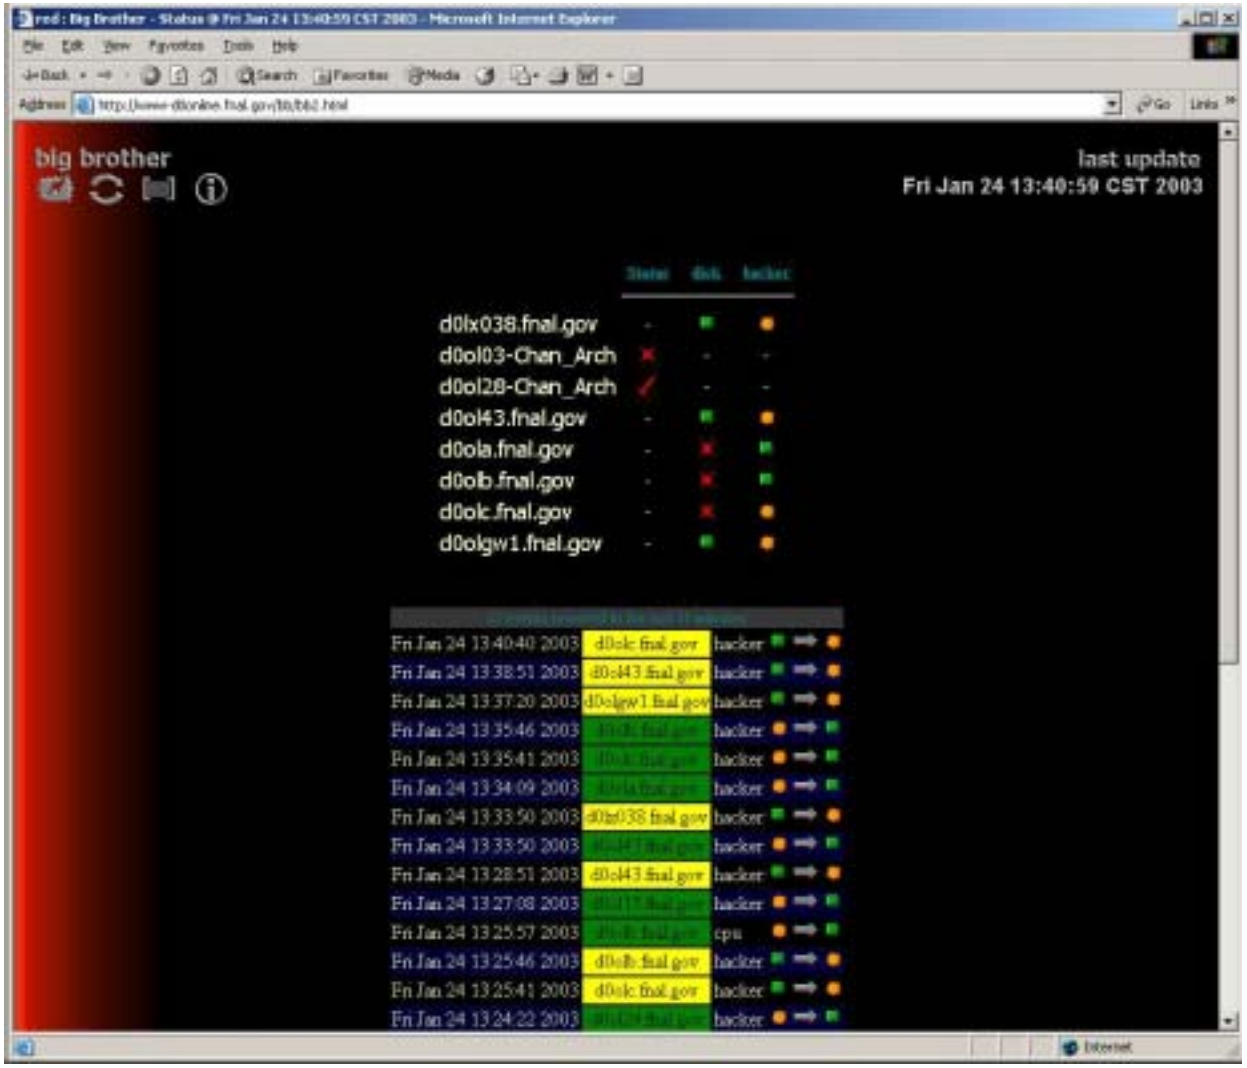

Summary display

click *button*

Monitoring – Big Brother

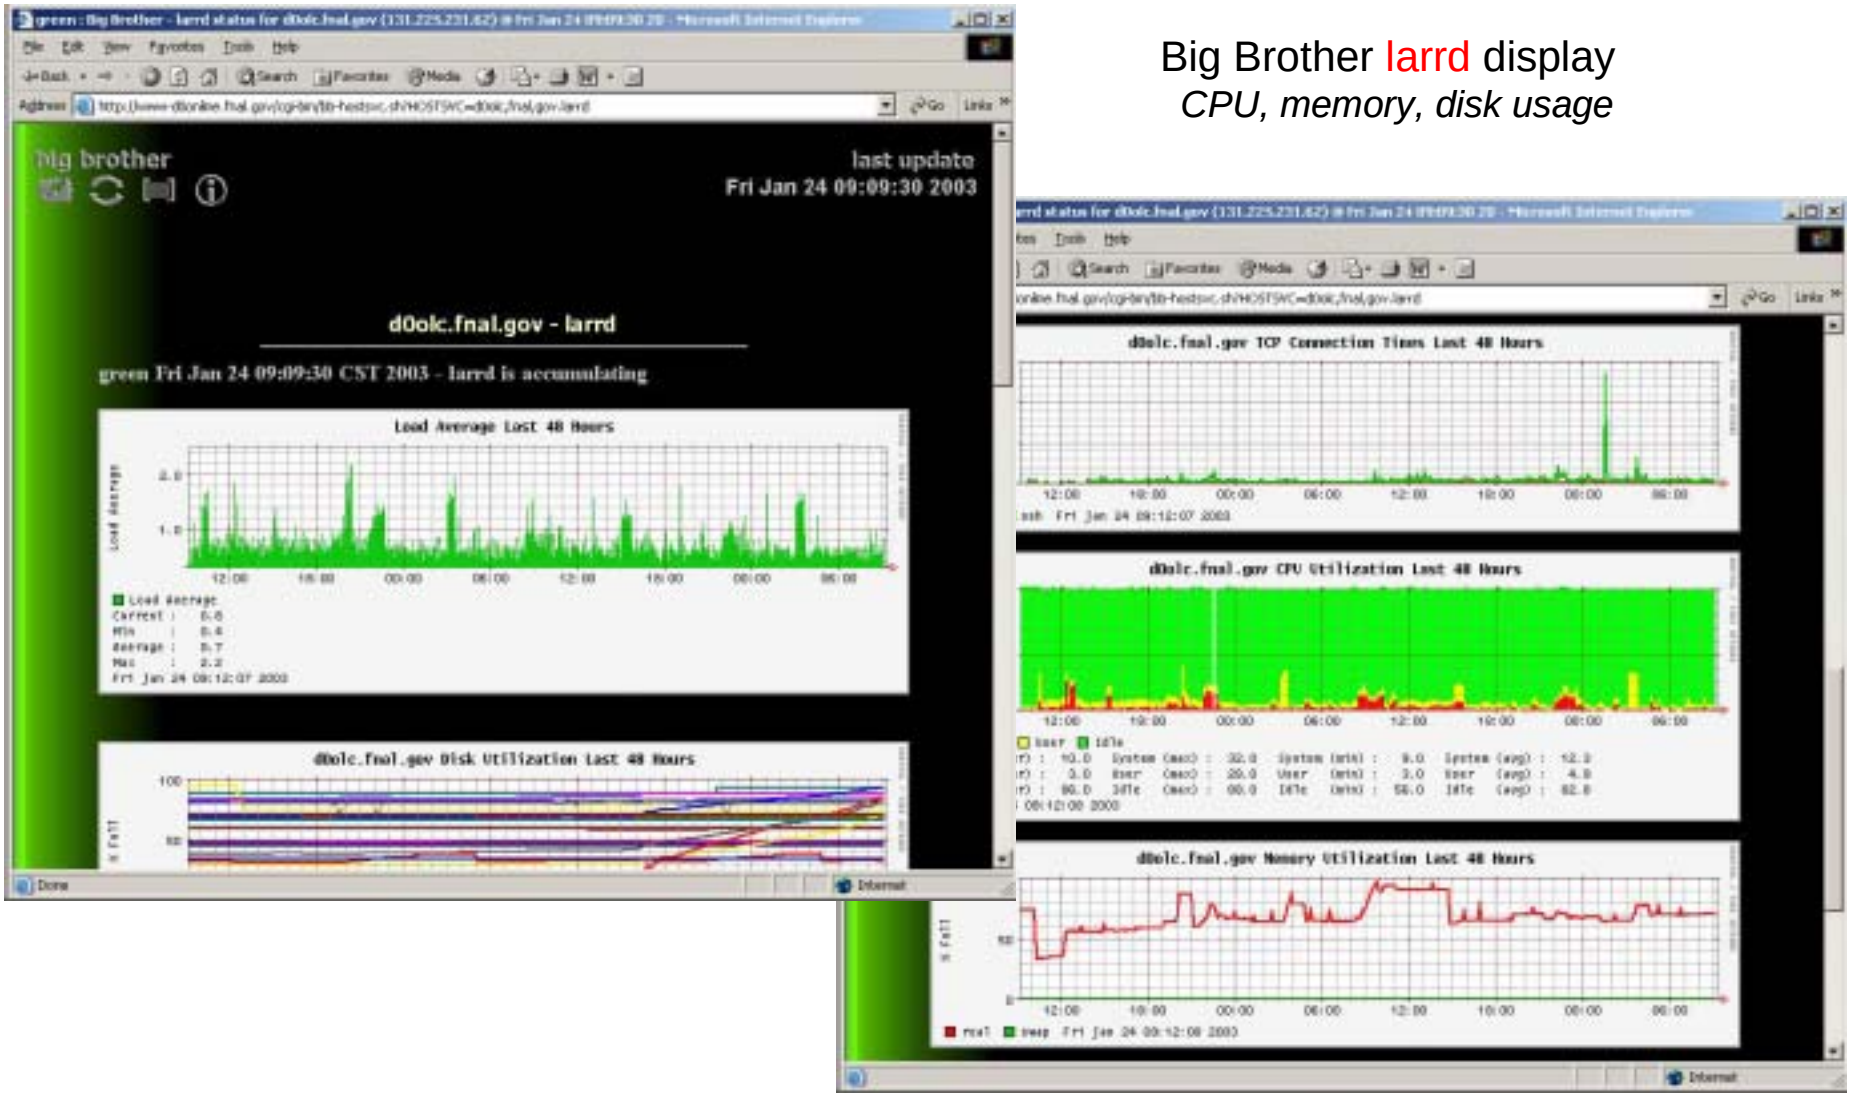

Big Brother larrd display

CPU, memory, disk usage

Monitoring – Big Brother

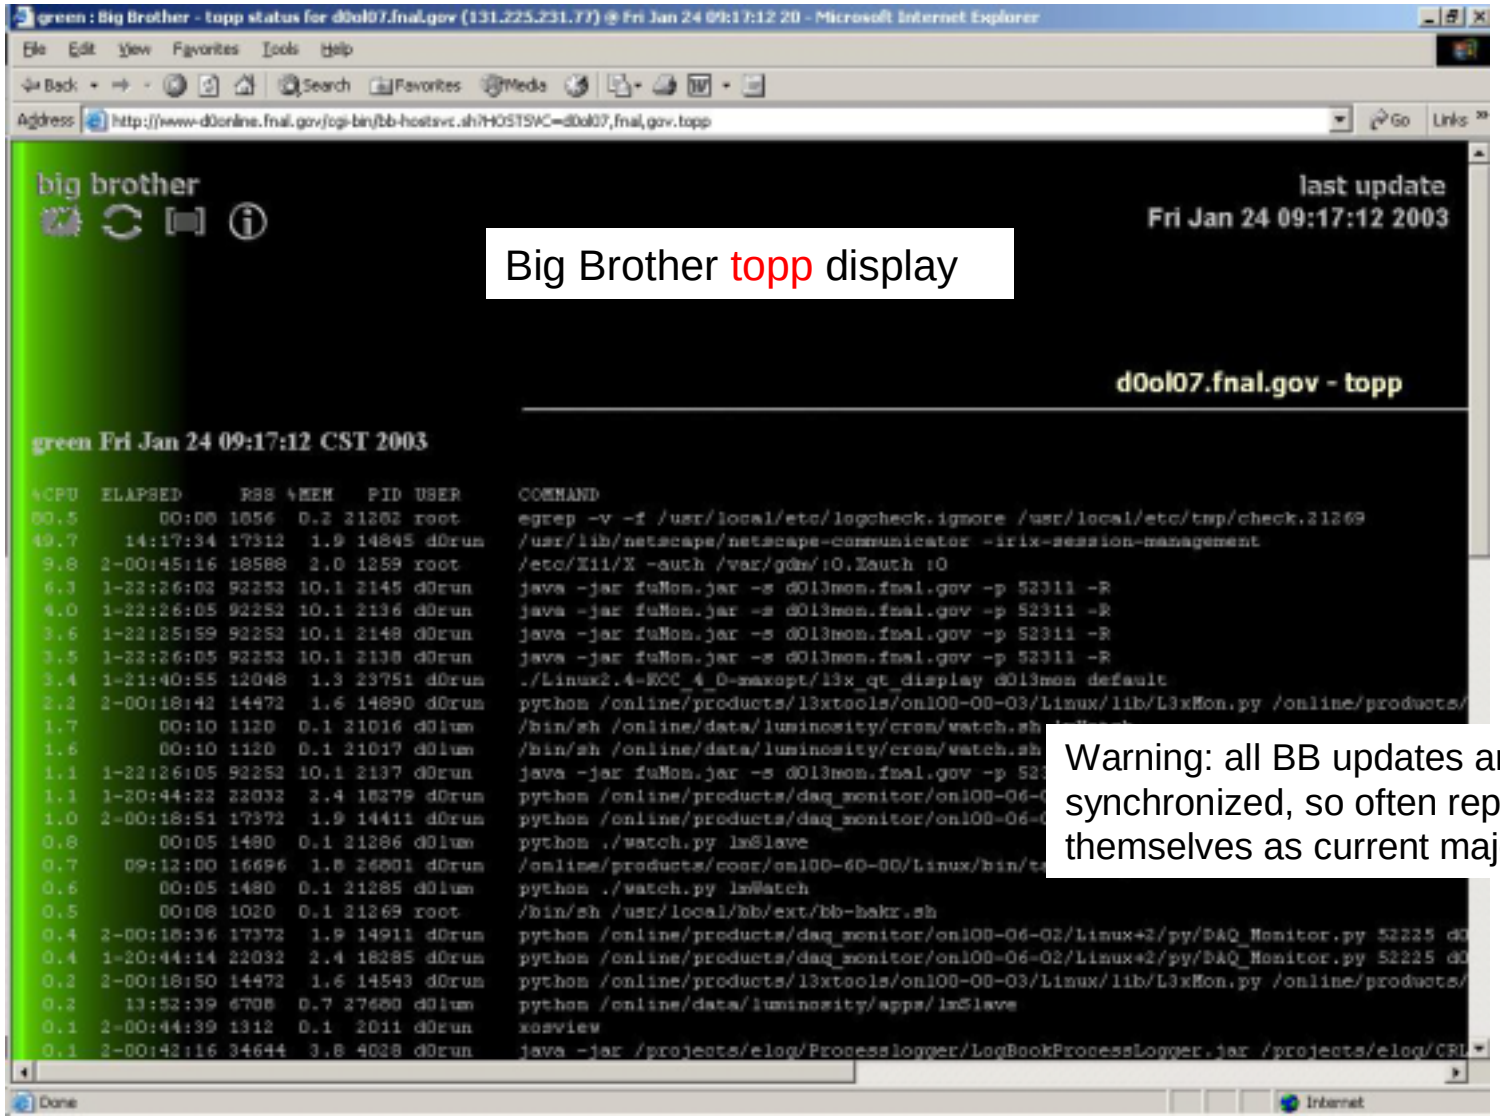

Monitoring – Big Brother

Big Brother disk display

Local disk usage

See d0ola/b/c for cluster disks

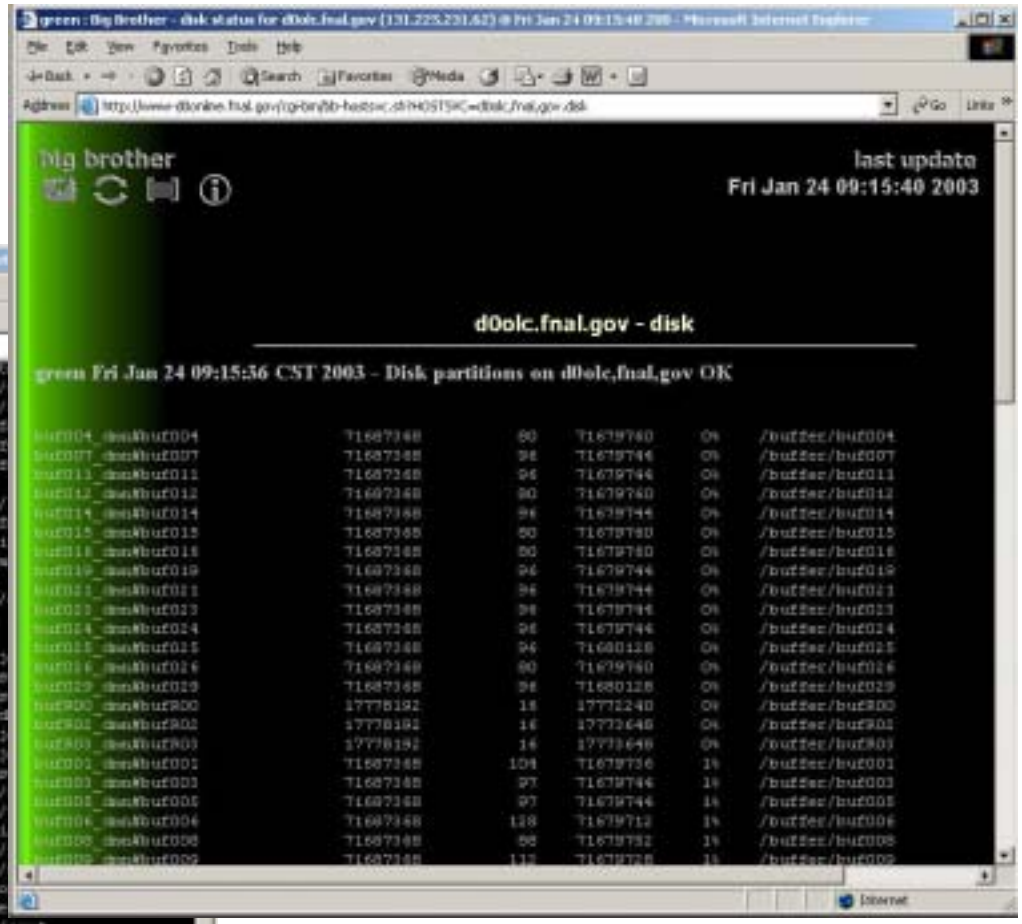

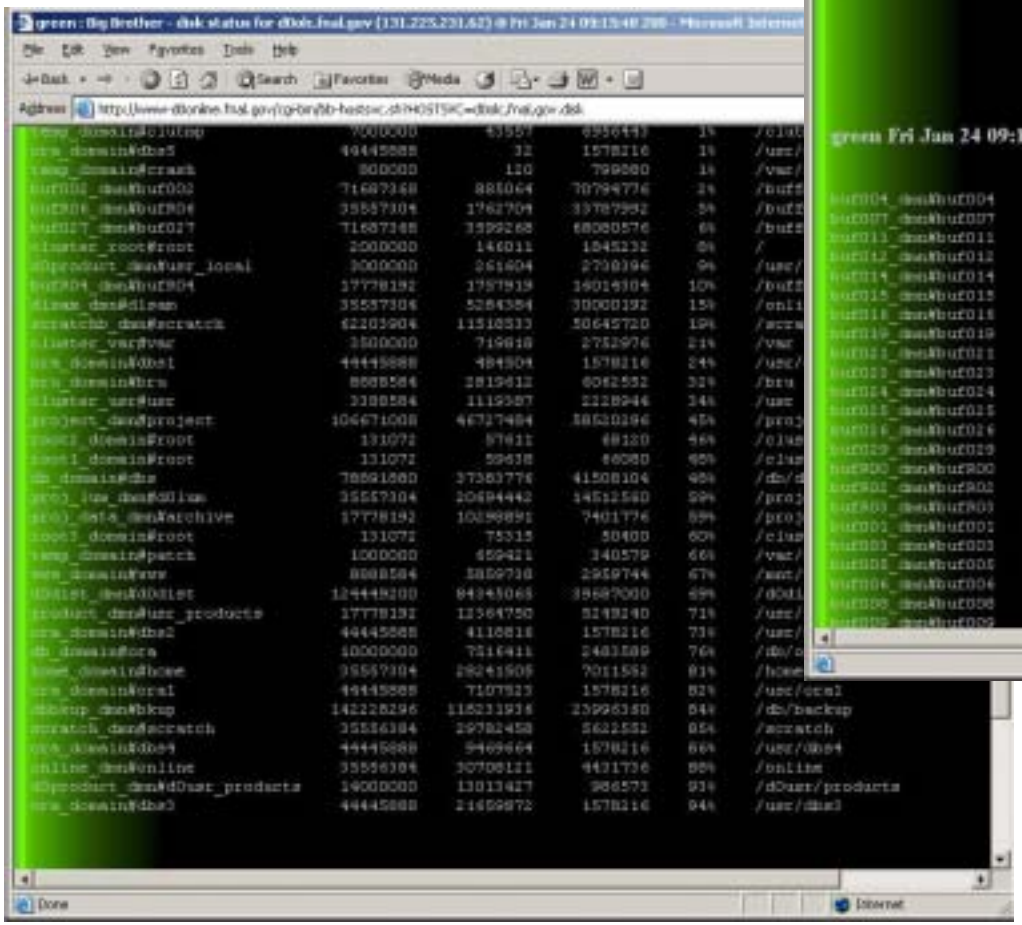

Monitoring – Big Brother

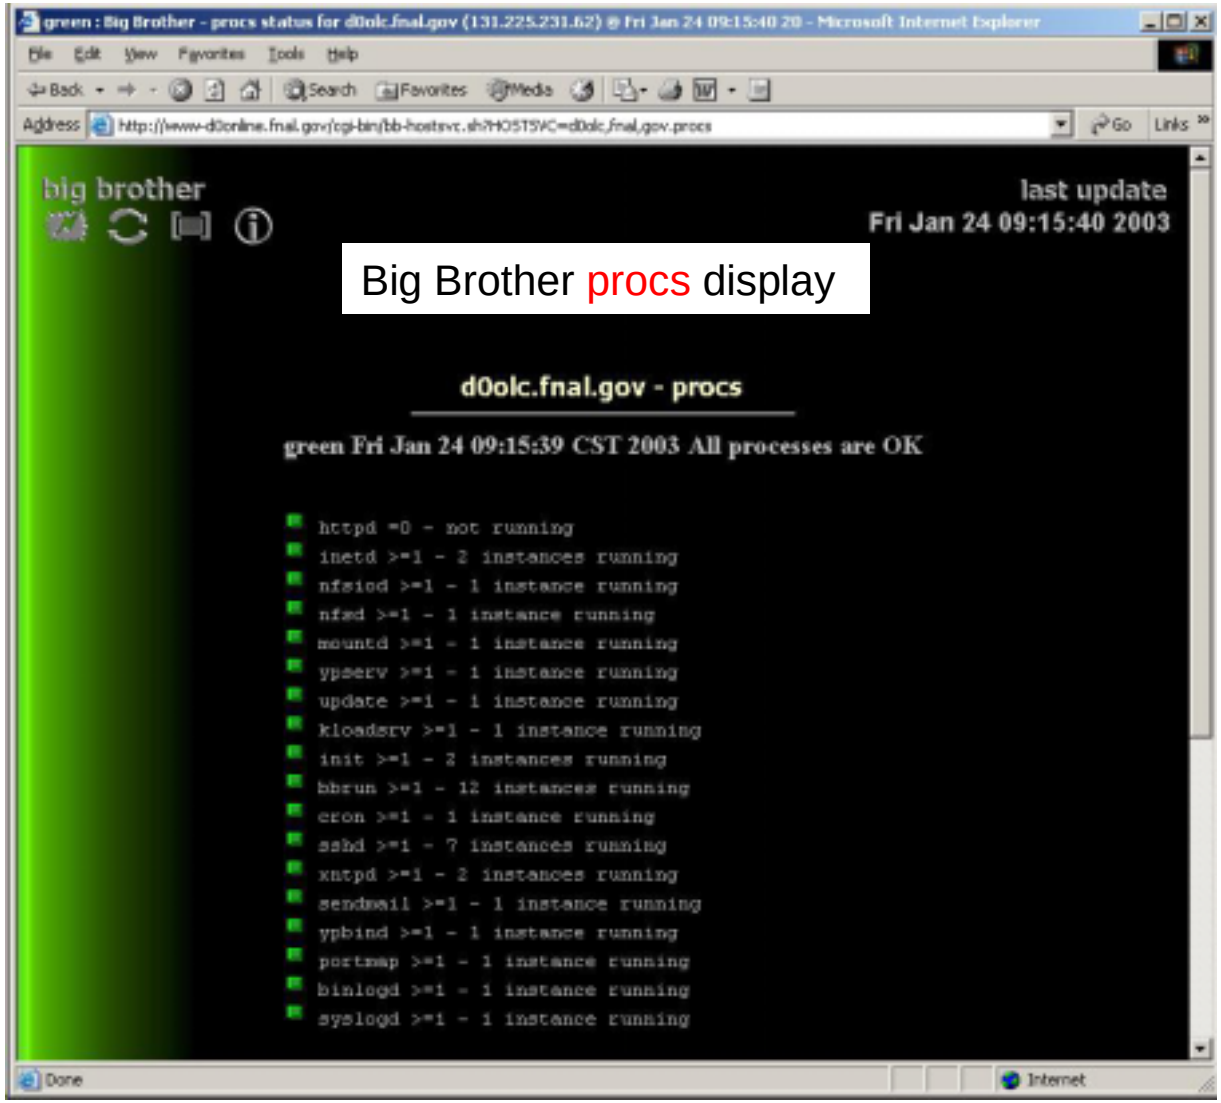

Web Servers

- There are several “internal” and “external” servers
  - Internal : visible only from within Online system
  - External : visible from anywhere
- One strategy is to mount / display from *same* disks
  - NFS mounted from a central server
  - Read-only mount to external servers
  - Appropriate ACL holes in router
  - Internal server:
    - http://www-d0ol.fnal.gov (alias for d0ol01)
  - External server:
    - http://www-d0online.fnal.gov (alias for d0online2)
- Other strategy is for server to act as client of internal node
  - Appropriate ACL holes in router
  - External server:
    - http://www-d0l3mon.fnal.gov

Control Room consoles

- Linux provides, by default, 6 serial and 1 graphical sessions
  - Graphical session is default
  - Switch among them with CTRL-ALT-F1 through CTRL-ALT-F7 keys
    - CTRL-ALT-F7 is the graphical session
- X is the windowing system for Linux

  Ref: http://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/x11.html
  - As opposed to Windows, where X has to be run “on top of” the native windowing system
  - The windowing system is the function of the “X server”
    - /etc/X11/X
      - Configured by /etc/X11/XF86config-4
        - Sets properties of graphics cards and monitors
      - Manages the DISPLAYs
      - Restart with CTRL-ALT-BACKSPACE – logs you out!
  - The X “display manager” runs to manage graphical logins
    - /usr/X11R6/bin/xdm
  - The X “window manager” runs upon login; we use fvwm
    - /usr/X11R6/lib/X11/fvwm2
      - Configured to set virtual windows, menus, etc
      - Restartable “hot” from menu

Useful tools

- Check usage
  - top
    - P (default) to sort by CPU usage
    - M to sort by memory usage
    - *Overutilization of swap space may indicate a memory leak problem*
- Files, sockets, etc
  - lsof
    - -p <pid>
    - -i [udp,tcp]:<port>
    - *Gives a lot of information on open files, network connections*
    - *(Tru64 nodes (d0ola/b/c) first require a 'setup lsof')*
- Check network connections:
  - netstat –a
- Check processes
  - ps [-lef]
  - ps –lfu <user name>
- Check memory utilization
  - vmstat [repeat period]
    - *Can see swap I/O, local disk I/O*

Common alarms

- Disk utilization
  - Questions to ask:
    - Is it a cluster or local disk?

      Shared mount on d0ola / d0olb / d0olc
    - Is it a critical disk?

      Cluster disks tend to be more important
    - Who's to blame?

      "User disks" (/home, /mnt/group)
    - How fast is it filling?

      Check periodically, or use "larrd" plots in Big Brother
- Memory utilization
  - Linux will try to utilize all the physical memory
    - Applications plus buffer cache
      - Use "free" to see the split
    - When physical memory exhausted, then swap space used
      - This can greatly degrade performance
  - Identify problematic applications